

Docket No. 50277-2249

Patent

UNITED STATES PATENT APPLICATION
FOR
DYNAMIC SELECTION OF FREQUENT ITEMSET COUNTING TECHNIQUE

INVENTORS:

WEI LI
JIANSHENG HUANG
ARI MOZES
SHIBY THOMAS
MARK DOUGLAS CALLAGHAN

ASSIGNEE:

ORACLE INTERNATIONAL CORPORATION
500 ORACLE PARKWAY
REDWOOD SHORES, CA 94065

PREPARED BY:

HICKMAN PALERMO TRUONG & BECKER LLP
1600 WILLOW STREET
SAN JOSE, CALIFORNIA 95125
(408) 414-1080

"Express Mail" mailing label number EV32335397645

Date of Deposit August 18, 2003

DYNAMIC SELECTION OF FREQUENT ITEMSET COUNTING TECHNIQUE

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is related to the following U.S. Patent Applications:

Serial No. 10/_____, entitled FREQUENT ITEMSET COUNTING USING CLUSTERED PREFIXES AND INDEX SUPPORT, (Attorney Docket No. 50277-2248), filed on the same day herewith; and

Serial No. 10/_____, entitled EXPRESSING FREQUENT ITEMSET COUNTING OPERATIONS, (Attorney Docket No. 50277-2250), filed on the same day herewith;

the contents of which are hereby incorporated by reference in their entirety.

FIELD OF THE INVENTION

[0002] The present invention relates to database operations and, more specifically, to techniques for performing frequent itemset operations.

BACKGROUND OF THE INVENTION

[0003] It is common for a data item that is stored in a database to have a logical relationship with other data items that are stored in the database. A set of data items that are related to each other is referred to herein as an "item group". An example of an item group is the set of all data items related to a particular employee (such as name, age, salary, etc.). Another example of an item group is the set of all data items that were purchased in a particular transaction (such as apples, bananas, and grapes).

[0004] A set of similar item groups is referred to herein as an "item group population". Relational database systems are frequently used to store information about large item group populations. For example, a relational database system may be used to store information

about all employees of a company. As another example, a relational database system may be used to store information about all sales transactions made at a given store, or at a large chain of stores.

[0005] Relational database systems are not only used to store information, but also to gather valuable intelligence based on the information that they store. For example, the management of a chain of stores may perform operations on the sales transaction information stored in a relational database to determine which stores are making the most sales, and which regions of the country are interested in particular products.

[0006] The most direct way to perform operations on data that is managed by a relational database server is to issue commands to the database server, where the commands specify the desired operations. In response to the commands, the relational database performs the desired operations and returns the results to the entity that issued the commands.

[0007] Of course, for the database server to execute the commands, the commands must conform to the database language that is supported by the database server. One database language that is supported by most relational database servers is SQL. Unfortunately, there is a limit to the type of operations that SQL directly supports. Operations that are not directly supported by SQL may be performed by specifying a series of SQL operations which, when executed in combination with each other, perform the desired unsupported operation.

[0008] Depending on the nature of the unsupported operation, the combination of SQL operations required to perform the unsupported operation may be quite complex. Further, amount of time and resources required to execute the series of operations may make the use of SQL impractical. Under these circumstances, it is often more efficient to simply export the data from the database and execute a software program specially designed to perform the desired operation on the expected data. If further operations are to be performed on the

results of the operation, then the results of the operation may be imported back into the database.

[0009] An example of a type of operation that, in general, cannot be performed efficiently using SQL operations is a frequent itemset operation. A frequent itemset operation is an operation that identifies which sets of items occur together most frequently in a particular item group population. For example, assume that a database stores information about sales transactions for a fruit market that sells apples, bananas and grapes. Assume further that ten percent of the sales transactions involve apples and bananas, that fifty percent of the sales transactions involve apples and grapes, and that ninety percent of the sales transactions involve grapes and bananas. If the frequent itemset operation uses a "frequency threshold" of seventy percent, then the results of the frequent itemset operation would include the itemset (grapes, bananas) but would exclude the itemsets (apples, grapes) and (apples, bananas). On the other hand, if the frequent itemset operation uses a frequency threshold of forty percent, then the results of the frequent itemset operation would include the itemsets (grapes, bananas) and (apples, grapes) but not the itemset (apples, bananas).

[0010] When performed using available SQL operations, frequent itemset operations typically require, among other things, so many join operations that performance is frequently unacceptable when the operation involves any sizable item group population. Consequently, for large item group populations, the technique of exporting the data from the database, performing the operation outside the database, and importing the results of the operation back into the database, may be the only acceptable alternative.

[0011] Based on the foregoing, it is desirable to provide a technique for performing frequent itemset operations in a manner that reduces the performance problems and the import/export hassle associated with the current techniques.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0013] FIG. 1 is a block diagram of a table that stores item groups according to a multiple-items-per-row format;

[0014] FIG. 2 is a block diagram of a table that stores item groups according to a single-item-per-row format;

[0015] FIG. 3 is a block diagram illustrating a set of single-item bitmaps, according to an embodiment of the invention;

[0016] FIG. 4 is a flowchart illustrating steps for managing bitmaps during a phase of a frequent itemset operation, according to an embodiment of the invention;

[0017] FIG. 5A is a block diagram of a prefix tree used for counting occurrences of combinations, according to an embodiment of the invention;

[0018] FIG. 5B is a block diagram of the prefix counting tree of FIG. 5A, where the counters associated with the leaf nodes have been incremented based on an item group population; and

[0019] FIG. 6 is a block diagram of a computer system upon which embodiments of the invention may be implemented.

DETAILED DESCRIPTION OF THE INVENTION

[0020] Techniques and systems are provided for performing frequent itemset operations. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

FUNCTIONAL OVERVIEW

[0021] Techniques are provided for (1) extending SQL to support direct invocation of frequent itemset operations, (2) improving the performance of frequent itemset operations by clustering itemset combinations to more efficiently use previously produced results, and (3) making on-the-fly selection of the occurrence counting technique to use during each phase of a multiple phase frequent itemset operation.

MULTIPLE-ITEMS-PER-ROW FORMAT ITEM GROUP POPULATIONS

[0022] Item group populations may be stored in any one of a variety of formats. According to one such format (referred to herein as the "multiple-items-per-row format"), the item group population is stored in a table, where

- each row of the table is associated with an item group, and
- each row of the table stores values that identify the members of the item group associated with the row.

[0023] An example of an item group population stored in the multiple-items-per-row format is illustrated in FIG. 1. Referring to FIG. 1, a table 100 stores data for an item group

population that includes three item groups. Each of the three rows (102, 104 and 106) of table 100 corresponds to an item group. For the purpose of explanation, the item groups associated with rows 102, 104 and 106 shall be referred to herein as item groups t1, t2 and t3, respectively.

SINGLE-ITEM-PER-ROW FORMAT ITEM GROUP POPULATIONS

[0024] Another format (referred to herein as the "single-item-per-row format") for storing an item group population involves storing the item group population in a two-column table, where for each row

- one column identifies an item group, and
- the other column identifies one of the items in the item group.

[0025] Since each row only identifies one item of an item group, an item group with N items will be represented by N rows. An example of an item group population stored in the single-item-per-row format is illustrated in FIG. 2. Referring to FIG. 2, table 200 includes data representing the same item group population as is represented in table 100 of FIG 1. However, in table 200, each row only identifies one (item group, item) association. Thus, item group t2, which contains three items, is represented in three rows.

OTHER FORMATS FOR STORING ITEM GROUP POPULATIONS

[0026] The multiple-items-per-row and single-item-per-row formats for storing item group populations, respectively illustrated in FIGS. 1 and 2, are merely two examples of a virtually unlimited number of ways item group population information may be represented. The techniques described herein are not limited to any particular representation format. However, regardless of the format used, information about the format is known by or

communicated to the database server to enable the database server to know how to process the information to perform frequent itemset operations.

[0027] An embodiment shall be described hereafter in which the database server understands both the multiple-items-per-row and single-item-per-row formats, and includes routines for receiving the item group information from each of the different formats. In embodiments that support multiple formats, the database statement that requests the performance of a frequent itemset operation may be used as the mechanism to communicate to the database server which input format to expect. Since the most widely used database language, SQL, does not have support for directly specifying a frequent itemset operation, nor for indicating an input format for the operation, extensions to the SQL syntax are described hereafter for allowing database commands to include this information.

SQL SUPPORT FOR FREQUENT ITEMSET OPERATIONS

[0028] According to one aspect of the invention, a system is provided that includes a database server capable of handling database commands that directly specify frequent itemset operations. According to one embodiment, the database server is configured to recognize, within a database statement that otherwise conforms to SQL, an SQL extension that specifies a frequent itemset operation. When executing the database statement, the database server performs the frequent itemset operation using an item group population specified in the SQL statement. According to one embodiment, the item group population may itself be a database statement (e.g. an SQL statement) of arbitrary complexity. In addition, the results produced by the frequent itemset operation may be further manipulated by the database statement.

[0029] For embodiments that support different input formats for the item group population, the SQL extension includes a mechanism for designating the specific format of

the input item group population. According to one embodiment, the database server includes direct support for frequent itemset operations, where

- SQL is extended to include table functions for directly invoking frequent itemset operations;
- support is provided for item group population input in both the single-item-per-row format and the multiple-items-per-row format; and
- the input item group population format is indicated by the name of the table function specified in the SQL statement.

[0030] The syntax, parameters and use of such table functions are explained in greater detail hereafter. However, those details merely provide an example of how SQL may be extended to provide direct support for frequent itemset operations. In alternative embodiments, frequent itemset operations may be directly specified using database language constructs other than table functions. Further, the specific syntax and parameters of the table functions may vary from implementation to implementation in embodiments that extend SQL to include table functions for direct frequent itemset operation support.

TABLE FUNCTION FOR SINGLE-ITEM-PER-ROW FORMAT INPUT

[0031] According to one embodiment, direct support for frequent itemset operations is provided by a database server configured to properly parse and execute SQL statements that support a table function named "FREQUENT ITEMSET_Transactional". The FREQUENT ITEMSET_Transactional function counts all frequent itemsets given a cursor for input data that is in the single-item-per-row format.

[0032] According to one embodiment, the table function for single-item-per-row format input has the following syntax:

```

FUNCTION FI_Transactional(
    tranx_cursor SYSREFCURSOR,
    support_threshold NUMBER,
    itemset_length_min NUMBER,
    itemset_length_max NUMBER,
    including_items SYSREFCURSOR,
    excluding_items SYSREFCURSOR)
RETURN TABLE OF ROW(
    itemset [Nested Table of Item Type DERIVED FROM tranx-cursor],
    support NUMBER,
    length NUMBER,
    total_tranx NUMBER);

```

[0033] The syntax example given above indicates support for a variety of parameters.

According to one embodiment, those parameters are used as follows:

[0034] tranx_cursor(IN): The cursor parameter that the user will supply when calling the function. It should return two columns in its returning row, the first column being an item group id, the second column being an item id.

[0035] support_threshold(IN): A fractional number. An itemset is called frequent if the number of transactions it occurs in divided by the total number of transactions exceeds the fraction. The parameter must be a NUMBER.

[0036] itemset_length_min(IN): Minimum length for interested frequent itemsets. The parameter must be a NUMBER.

[0037] Itemset_length_max(IN): Maximum length for interested frequent itemsets. This parameter must be a NUMBER.

[0038] including_items (IN): A cursor from which a list of items can be fetched. At least one item from the list should appear in frequent itemsets that will be returned.

[0039] excluding_items(IN): A cursor from which a list of items can be fetched. No item from the list can appear in frequent itemsets that will be returned.

[0040] support(OUT): number of transactions in which a frequent itemset occurs. This will be returned as a NUMBER.

[0041] itemset(OUT): A collection of items which is computed to be frequent. This will be returned as a nested table of an item type which is the item column type of the input cursor.

[0042] length(OUT): Number of items in a frequent itemset. This will be returned as a NUMBER.

[0043] total_tranx(OUT): Total transaction count. This will be returned as a NUMBER.

[0044] The FI_Transactional function's behavior is illustrated with the following example. Suppose that the input table tranx_table_in contains the data illustrated in table 200 of FIG. 2. Assume further that the user is trying to find itemsets that satisfy a support-threshold of 60% and have an itemset-length greater than 1 (namely, (apple, banana)).

[0045] itemset=(apple, banana), support=2, length=2, total_tranx=3

[0046] The following illustrates how the FI_Transactional function may be used in an otherwise standard SQL query to satisfy the user request:

```
select itemset, support, length, total_tranx
from table(sys.dbms_frequent_itemset.fi_transactional(
    cursor(select tid, iid from tranx_table_in),
    0.6, 2, 5,
    NULL, NULL));
```

[0047] Applications can predefine a nested table type of the given input item type and cast itemset into the predefined type before further processing (e.g. loading into a table).

[0048] The following example illustrates how to include certain items and exclude certain items in the counting.

```
select itemset, support, length, total_tranx
from table(sys.dbms_frequent_itemset.fi_transactional(
    cursor(select tid, iid from tranx_table_in_IF1),
    0.6, 2, 5,
    cursor(select * from table(sys.ora_fi_char_nt(apple, banana))),
    cursor(select * from table(sys.ora_fi_char_nt(milk, orange))));
```

[0049] Using the including/excluding items parameter, the execution of the statement can be optimized by ignoring itemsets that are not desired by the application.

TABLE FUNCTION FOR MULTIPLE-ITEMS-PER-ROW FORMAT INPUT

[0050] According to one embodiment, direct support for frequent itemset operations is provided by a database server configured to properly parse and execute SQL statements that support a table function named "FI_Horizontal". The FI_Horizontal function counts all frequent itemsets given a cursor for input data that is in the multiple-items-per-row format. According to one embodiment, the syntax for the FI_Horizontal function is:

```
FUNCTION FI_Horizontal(
    tranx-cursor SYSREFCURSOR,
    support-threshold NUMBER,
    itemset-length-min NUMBER,
    itemset-length-max NUMBER,
```

```

including-items SYSREFCURSOR,
excluding-items SYSREFCURSOR)
RETURN TABLE OF ROW(
itemset [Nested Table of Item Type DERIVED FROM tranx-cursor],
support NUMBER,
length NUMBER,
total_tranx NUMBER);

```

[0051] The syntax example given above indicates support for a variety of parameters.

According to one embodiment, those parameters are used as follows:

[0052] `tranx_cursor(IN)`: The cursor parameter that the user will supply when calling the function. Each non-null column of the cursor represents a an item. All columns of the cursor have the same data type.

[0053] `support_threshold(IN)`: A fractional number. An itemset is called frequent if the number of transactions it occurs in divided by the total number of transactions exceeds the fraction. The parameter must be a NUMBER.

[0054] `itemset_length_min(IN)`: Minimum length for interested frequent itemsets. The parameter must be a NUMBER.

[0055] `Itemset_length_max(IN)`: Maximum length for interested frequent itemsets. This parameter must be a NUMBER.

[0056] `including_items (IN)`: A cursor from which a list of items can be fetched. At least one item from the list should appear in frequent itemsets that will be returned.

[0057] `excluding_items(IN)`: A cursor from which a list of items can be fetched. No item from the list can appear in frequent itemsets that will be returned.

[0058] support(OUT): number of transactions in which a frequent itemset occurs. This will be returned as a NUMBER.

[0059] itemset(OUT): A collection of items which is computed to be frequent. This will be returned as a nested table of an item type which is the item column type of the input cursor.

[0060] length(OUT): Number of items in a frequent itemset. This will be returned as a NUMBER.

[0061] total_tranx(OUT): Total transaction count. This will be returned as a NUMBER.

[0062] The following example illustrates how to use this function for data that is stored in the 'multiple-items-per-row' format. For the purpose of illustration, assume that a table `horiz_table_in` contains the columns and data illustrated in table 100 of FIG. 1. Further assume that a user wishes to find out what combinations of items are frequent with a given support threshold of 30%, requiring itemsets containing at least one of (apple, banana), but excluding any of (milk, orange) in any itemset. This information can be obtained by issuing the following SQL query:

```
select itemset, support, length, total_tranx
from table(sys.dbms_frequent_itemset.fi_horizontal(
    cursor(select iid1, iid2, iid3, iid4, iid5
        from horiz_table_in),
        0.3, 2, 5,
        cursor(select * from table(sys.ora_fi_number_nt(apple, banana))),
        cursor(select * from table(sys.ora_fi_number_nt(milk, orange))))));
```

[0063] The benefit of this table function is that if the application already has data in multiple-items-per-row format, the database can skip the step of transforming rows from the single-item-per-row format into the multiple-items-per-row format.

PERFORMING FREQUENT ITEMSET OPERATIONS

[0064] When a database server receives a database statement, the database parses the database statement to determine the operations that the statement requires. To execute the statement, the database server must then perform those operations. Thus, for a database server to be able to handle statements that specify frequent itemset operations, the database server must be able to perform frequent itemset operations, and not merely be able to detect when a database statement specifies a frequent itemset operation.

[0065] There are many techniques that a database server may use to perform frequent itemset operations. However, not all techniques can be performed as efficiently as others. As shall be described in greater detail hereafter, a technique is provided for arranging the order in which itemsets are processed during a frequent itemset operation, such that the itemsets are clustered in a manner that makes more efficient use of available resources.

MULTI-PHASE CANDIDATE ELIMINATION

[0066] During a frequent itemset operation, the database server determines which of the frequent itemset candidates qualify as frequent itemsets, and which ones are disqualified for failing to satisfy a condition of the frequent itemset operation. At the start of a frequent itemset operation, every possible combination of items is considered to be a frequent itemset candidate. Candidates are then eliminated, as it is determined that they fail to satisfy a condition of the frequent itemset operation.

[0067] According to one embodiment, the candidate elimination process is performed in phases, where each phase focuses on itemset candidates that have a particular number of items. For example, the first phase focuses on itemset candidates that have one item. The second phase focuses on itemset candidates that have two items. The elimination process continues until no unclassified candidates remain, or until all itemset candidates that have M items have been processed, where M is the maximum itemset size specified in the database statement.

[0068] According to one embodiment, during the Nth phase, not every possible combination of N items is considered as a candidate. Rather, the results of the N-1 phase are used to compute which N-item combinations are candidates. For example, assume that, during the third phase it is determined that the combination (a, b, c) is not a frequent itemset because it does not satisfy the frequency criteria. If (a, b, c) does not satisfy the frequency criteria, then it is not possible for any four-item combination that includes (a, b, c) to satisfy the frequency criteria. Consequently, the initial set of candidates for the Nth phase are generated based on only the combinations that satisfied the frequency criteria during the N-1 phase.

[0069] According to one embodiment, an N-item combination is only considered as a candidate if all of its N-1 combination satisfy the frequency criteria.

USING BITMAPS TO IDENTIFY FREQUENT ITEMSETS

[0070] A candidate combination is not a resultant frequent itemset if (1) the candidate does not occur as frequently as required by the specified frequency threshold, or (2) the candidate fails to meet some other condition specified for the operation (e.g. the candidate has fewer members than the specified minimum, or the candidate does not include a

particular item that must be included). One technique to determine whether a particular candidate satisfies a frequency threshold involves

- (1) creating a bitmap for each possible item,
- (2) performing a Boolean AND operation between the single-item bitmaps to create a bitmap that represents the candidate combination,
- (3) determining how many item groups include the combination based on the bitmap of the combination, and
- (4) determining whether the candidate combination satisfies the frequency threshold based on the number of item groups that include the combination.

[0071] FIG. 3 is a diagram of a table of single item bitmaps 300. In the illustrated table, each item has a corresponding column, and each item group has a corresponding row. The intersection of a row and column of the table constitutes a cell. The value stored in a given cell indicates whether the item group associated with the row includes the specific item associated with the column.

[0072] For example, rows 302, 304 and 306 are respectively associated with item groups t1, t2 and t3. The columns of the table are associated with items "apple", "banana", "milk", and "orange". Cell 308, which is at the intersection of row 306 and column "orange" contains a "1" to indicate that the item "orange" is included in the transaction t3 associated with row 306.

[0073] The string of values in the column associated with an item constitutes the bitmap for that item. For example, the apple column of rows 302, 304 and 306 respectively store the values "1", "1" and "0". Consequently, the bitmap for "apple" is "110". Similarly, the bitmap for "banana" is "110", the bitmap for "milk" is "010" and the bitmap for "orange" is "001".

[0074] For the purpose of explanation, the notation B(a, b, c) will be used herein to indicate the bitmap for the combination that consists of items a, b, and c. Thus, based on the table illustrated in FIG. 3, B(apple) is "110" and B(orange) is "001".

[0075] The bitmap for a multi-item combination may be generated by performing an AND operation between the bitmaps for the individual items in the combination. For example, B(apple, milk) may be generated by performing an AND between the B(apple) "110" and B(milk) "010". In this example, the resulting bitmap B(apple, milk) is "010". Similarly, the bitmaps for one or more smaller combinations may be ANDed together to produce the bitmap for a larger combination. For example, the bitmap B(apple, milk) "010" may be ANDed with the bitmap B(banana) "001" to produce the bitmap B(apple, milk, banana) "000".

[0076] Each value in a bitmap for a combination indicates whether the combination exists (is "TRUE") or does not exist (is "FALSE") in a given item group. Therefore, the total number of TRUE indications in a bitmap indicates how many item groups in the item group population include the combination associated with the bitmap. In the table illustrated in FIG. 3, the value "1" is used to indicate TRUE. Thus, the bitmap "010" for the combination (apple, milk) indicates that only one item group in the relevant population has the combination (apple, milk). Similarly, the bitmap "000" for the combination (apple, milk, banana) indicates that no item group in the relevant population has the combination (apple, milk, banana).

BITMAP PROLIFERATION

[0077] Because the bitmaps of smaller combinations can be combined to create the bitmaps for larger combinations, various techniques are described herein in which the phases

of a frequent itemset operation begin with a single-item phase and proceed to phases for progressively larger combinations. For example, during the first phase, bitmaps are made for single-item combinations. During the second phase, the single-item bitmaps are used to make bitmaps for two-item combinations. During the third phase, the two-item bitmaps are combined to make bitmaps for three-item combinations.

[0078] Unfortunately, due to the amount of resources that retention of the bitmaps would require, it is rarely efficient to retain all bitmaps from previously-processed combinations. This is particularly true for item group populations that include large numbers of item groups (thereby increasing the size of each bitmap) and large numbers of possible items (thereby greatly increasing the number of combinations).

CLUSTERING COMBINATIONS

[0079] As explained above, the bitmap of a multiple-item combination may be generated by combining two or more lower-ordered bitmaps. The lower-ordered bitmaps that can be used to generate the bitmap of a multiple-item combination are referred to herein as the "base bitmaps" of the multiple-item combination. According to one embodiment, the problem of bitmap proliferation is reduced by (1) clustering the candidate combinations that are to be processed during a phase, and (2) selectively discarding bitmaps after processing the cluster(s) to which the bitmaps apply.

[0080] As used herein, the term "cluster" refers to a set of combinations that have a base bitmap in common. The base bitmap that all combinations in a particular cluster have in common is referred to herein as the "cluster bitmap" for that cluster. For example, consider the combinations (a, b, c, d), (a, b, c, e), (a, b, c, f), (a, m, n, o), (a, m, n, p), (a, m, n, q). These combinations can be separated into a first cluster that includes (a, b, c, d), (a, b, c, e),

(a, b, c, f), and a second cluster that includes (a, m, n, o), (a, m, n, p), (a, m, n, q). In this example, all combinations in the first cluster include the combination (a, b, c). Thus, the bitmap B(a, b, c) is the cluster bitmap for the first cluster. Similarly, all combinations in the second cluster include the combination (a, m, n). Thus, the bitmap B(a, m, n) is the cluster bitmap for the second cluster.

[0081] As shall be illustrated in greater detail hereafter, the concept of clustering is multi-layered. For example, all combinations that include the prefix (a, b, c, ...) may be considered a first cluster, and all combinations that include the prefix (a, b, ...) may be considered a second cluster, where the second cluster is a superset of the first cluster.

[0082] By processing combinations in a sequence that is based on clusters, it is possible to know when cluster bitmaps can be discarded. For example, the cluster bitmap B(a, b, c) may be discarded after being used to generate the bitmaps for the combinations that belong to the first cluster. Similarly, the cluster bitmap B(a, m, n) may be discarded after being used to generate the bitmaps for the combinations that belong to the second cluster.

[0083] Various techniques may be used to cluster the combinations that are being processed at a particular phase of a frequent itemset operation. For example, according to one embodiment, the items within a combination are sorted based on some ordering criteria, and then the combinations themselves sorted relative to each other based on the same ordering criteria. For example, assume that the combinations are (m, a, n, q), (d, b, c, a), (o, m, n, a), (a, e, c, b), (c, b, a, f), (a, m, n, p). Further assume that the sorting is performed based on alphabetic order. Under these circumstances, sorting the items within each combination would result in the combinations (a, m, n, q), (a, b, c, d), (a, m, n, o), (a, b, c, e), (a, b, c, f), (a, m, n, p). After the items within each combination are sorted, the combinations themselves are sorted relative to each other, resulting in the sequence: (a, b, c, d), (a, b, c, e),

(a, b, c, f), (a, m, n, o), (a, m, n, p), (a, m, n, q). Significantly, in the resulting sequence, the combinations with the base bitmap B(a, b, c) are clustered together, and the combinations with the base bitmap B(a, m, n) are clustered together.

[0084] Another technique for clustering combinations involves hashing the combinations into buckets based on sub-combinations. For example, all combinations that include items a, b and c may be hashed into one bucket, while all combinations that include items a, m and n may be hashed into another bucket. All of the combinations that fall into a particular bucket constitute a cluster, where the bitmaps of the combination used to establish the bucket are the cluster bitmaps for the cluster. For example, if all combinations that include a, b and c are hashed to a bucket X, then the bitmap B(a, b, c) would be the cluster bitmap for the cluster that consists of the combinations that have been hashed to bucket X.

TEMPORARILY RETAINING CLUSTER BITMAPS

[0085] According to one embodiment, only single-item bitmaps are retained between phases. Cluster bitmaps, on the other hand, are created prior to processing the clusters with which they are associated, and discarded after processing all combinations in the clusters with which they are associated. According to one embodiment, all other bitmaps are discarded as soon as they have been used to determine whether the corresponding combination satisfies the frequency criteria.

[0086] For example, assume that (a, b, c, d), (a, b, c, e), (a, b, c, f), (a, m, n, o), (a, m, n, p), and (a, m, n, q) are the candidates at the outset of the four-item combination phase. Initially, the only pre-existing bitmaps are the single-item bitmaps. In this example, the first cluster, which includes the combinations (a, b, c, d), (a, b, c, e), (a, b, c, f), is associated with the cluster bitmap B(a, b, c). Thus, prior to processing the first cluster, the bitmap B(a, b, c)

is generated from the single-item bitmaps, and then used along with the single-item bitmaps to generate the bitmaps for each of the combinations in the first cluster.

[0087] When a bitmap is generated for a combination in the first cluster, the bitmap is used to determine whether the combination satisfies the frequency threshold, and is then discarded. After all of the combinations in the first cluster have been evaluated, the bitmap $B(a, b, c)$ is also discarded.

[0088] At this point, the cluster bitmap $B(a, m, n)$ is created for the second cluster, and then used along with the single-item bitmaps to generate the bitmaps for each of the combinations in the second cluster. After all of the combinations in the second cluster have been evaluated, the bitmap $B(a, m, n)$ is discarded.

RETAINING BASE BITMAPS

[0089] In the example given above, cluster bitmaps are retained until all combinations in their corresponding cluster have been evaluated. According to one embodiment, all intermediary bitmaps that are generated to evaluate a given combination are also retained, as long as such intermediary bitmaps continue to be base bitmaps of the combination that is being evaluated.

[0090] FIG. 4 is a flowchart showing steps for evaluating combinations, according to an embodiment where intermediary bitmaps are retained as long as they are base bitmaps of the combination being evaluated. For the purpose of explanation, it shall be assumed that (a, b, c, d) , (a, b, c, e) , (a, b, c, f) , (a, b, n, o) , (a, m, n, p) , and (a, m, n, q) are the candidates at the outset of the four-item combination phase. At step 400, the appropriate single-item bitmap is selected for evaluating the first combination (a, b, c, d) . Since the first combination is (a, b, c, d) , the single-item bitmap $B(a)$ will initially be selected.

[0091] In step 402, the intermediary bitmaps necessary to evaluate the current combination are generated. In the present example, the intermediary bitmaps that are generated to evaluate the combination (a, b, c, d) include B(a, b) and B(a, b, c). In step 404, the current combination (a, b, c, d) is evaluated by combining the largest retained bitmap B(a, b, c) with the appropriate single value bitmap B(d). After (a, b, c, d) has been evaluated, it is determined whether there are any more combinations to be processed (step 406). If not, then the current phase is done (step 408).

[0092] In the present example, more combinations need to be processed, so control passes to step 410. At step 410, the next combination (a, b, c, e) is selected. At step 412, it is determined whether the largest retained bitmap B(a, b, c) is a base bitmap of the current combination (a, b, c, e). In the present example, the largest retained bitmap B(a, b, c) is a base bitmap of (a, b, c, e). Therefore, no retained bitmap is discarded, and control passes to step 402. In step 402, any necessary intermediary bitmaps are generated to evaluate (a, b, c, e). In the present example, the largest retained bitmap is only one item smaller than the combination being evaluated, so no intermediary bitmaps need to be generated. At step 404, the largest retained bitmap B(a, b, c) is used to evaluate the combination (a, b, c, e).

[0093] After (a, b, c, e) has been evaluated, it is determined whether there are any more combinations to be processed (step 406). If not, then the current phase is done (step 408).

[0094] In the present example, more combinations need to be processed, so control passes to step 410. At step 410, the next combination (a, b, c, f) is selected. At step 412, it is determined whether the largest retained bitmap B(a, b, c) is a base bitmap of the current combination (a, b, c, f). In the present example, the largest retained bitmap B(a, b, c) is a base bitmap of (a, b, c, f). Therefore, no retained bitmap is discarded, and control passes to step 402. In step 402, any necessary intermediary bitmaps are generated to evaluate (a, b, c,

f). In the present example, the largest retained bitmap is only one item smaller than the combination being evaluated, so no intermediary bitmaps need to be generated. At step 404, the largest retained bitmap B(a, b, c) is used to evaluate the combination (a, b, c, f).

[0095] After (a, b, c, f) has been evaluated, it is determined whether there are any more combinations to be processed (step 406). If not, then the current phase is done (step 408).

[0096] In the present example, more combinations need to be processed, so control passes to step 410. At step 410, the next combination (a, b, n, o) is selected. At step 412, it is determined whether the largest retained bitmap B(a, b, c) is a base bitmap of the current combination (a, b, n, o). In the present example, the largest retained bitmap B(a, b, c) is not a base bitmap of (a, b, n, o). Therefore, control passes to step 414.

[0097] In step 414, it is determined whether the largest retained bitmap is a single-item bitmap. In the present example, the largest retained bitmap B(a, b, c) is not a single-item bitmap, so control passes to step 416 and the largest retained bitmap is discarded. After discarding the bitmap B(a, b, c), the largest retained bitmap is the bitmap B(a, b). At step 412, it is determined whether the largest retained bitmap B(a, b) is a base bitmap of (a, b, n, o). In the present example, B(a, b) is a base bitmap of (a, b, n, o), so control passes to step 402.

[0098] In step 402, any necessary intermediary bitmaps are generated to evaluate (a, b, n, o). In the present example, the bitmap B(a, b, n) is generated. Consequently, the largest retained bitmap is the bitmap B(a, b, n). At step 404, the largest retained bitmap B(a, b, n) is used to evaluate the combination (a, b, n, o).

[0099] After (a, b, n, o) has been evaluated, it is determined whether there are any more combinations to be processed (step 406). If not, then the current phase is done (step 408).

[0100] In the present example, more combinations need to be processed, so control passes to step 410. At step 410, the next combination (a, m, n, p) is selected. At step 412, it is determined whether the largest retained bitmap B(a, b, n) is a base bitmap of the current combination (a, m, n, p). In the present example, the largest retained bitmap B(a, b, n) is not a base bitmap of (a, m, n, p). Therefore, control passes to step 414.

[0101] In step 414, it is determined whether the largest retained bitmap is a single-item bitmap. In the present example, the largest retained bitmap B(a, b, n) is not a single-item bitmap, so control passes to step 416 and the largest retained bitmap is discarded. After discarding the bitmap B(a, b, c), the largest retained bitmap is the bitmap B(a, b).

[0102] At step 412, it is determined whether the largest retained bitmap B(a, b) is a base bitmap of the current combination (a, m, n, p). In the present example, the largest retained bitmap B(a, b) is not a base bitmap of (a, m, n, p). Therefore, control passes to step 414.

[0103] In step 414, it is determined whether the largest retained bitmap is a single-item bitmap. In the present example, the largest retained bitmap B(a, b) is not a single-item bitmap, so control passes to step 416 and the largest retained bitmap is discarded. After discarding the bitmap B(a, b), the largest retained bitmap is the bitmap B(a).

[0104] At step 412, it is determined whether the largest retained bitmap B(a) is a base bitmap of the current combination (a, m, n, p). In the present example, the largest retained bitmap B(a) is a base bitmap of (a, m, n, p). Therefore, no retained bitmap is discarded, and control passes to step 402. In step 402, any necessary intermediary bitmaps are generated to evaluate (a, m, n, p). In the present example, the intermediary bitmaps B(a, m) and B(a, m, n) are generated. At step 404, the largest retained bitmap B(a, m, n) is used to evaluate the combination (a, m, n, p).

[0105] After (a, m, n, p) has been evaluated, it is determined whether there are any more combinations to be processed (step 406). If not, then the current phase is done (step 408).

[0106] In the present example, more combinations need to be processed, so control passes to step 410. At step 410, the next combination (a, m, n, q) is selected. At step 412, it is determined whether the largest retained bitmap B(a, m, n) is a base bitmap of the current combination (a, m, n, q). In the present example, the largest retained bitmap B(a, m, n) is a base bitmap of (a, m, n, q). Therefore, no retained bitmap is discarded, and control passes to step 402. In step 402, any necessary intermediary bitmaps are generated to evaluate (a, m, n, q). In the present example, the largest retained bitmap is only one item smaller than the combination being evaluated, so no intermediary bitmaps need to be generated. At step 404, the largest retained bitmap B(a, m, n) is used to evaluate the combination (a, m, n, q).

[0107] After (a, m, n, q) has been evaluated, it is determined whether there are any more combinations to be processed (step 406). Because (a, m, n, q) was the last 4-item candidate, the current phase is done (step 408). According to one embodiment, at the end of each phase, all of the retained bitmaps are discarded, except the single-item bitmaps.

STORING BITMAPS IN INDEXES

[0108] In most computing environments, volatile memory is smaller but faster than non-volatile memory. Thus, to improve the speed of frequent itemset operations, all of the bitmaps that are generated during a frequent itemset operation may be maintained in volatile memory. However, even large systems may not have enough volatile memory to maintain all data involved in a frequent itemset operation in volatile memory.

[0109] According to one embodiment, the results of each phase are written to non-volatile memory. During a subsequent phase, the results of a previous phase are loaded from

non-volatile memory into volatile memory as those results become needed in the subsequent phase.

[0110] When non-volatile memory is used for bitmap retention, some or all of the bitmaps that are generated during one phase of a frequent itemset operation may be stored in an index in non-volatile memory. According to one embodiment, the bitmaps of all N-item combinations that satisfied the frequency criteria are stored in a non-volatile index so that they can be later used to create the bitmaps for N+1-item combinations.

[0111] The combination associated with a bitmap may be used as the index key for determining where within the index to place the entry for the bitmap. According to one embodiment, a bitmap index is used to store the bitmaps. Each entry within the bitmap index is associated with a combination, and includes the bitmap for the combination. Within the bitmap index, the bitmaps may be compressed to preserve space. In addition, when a stored bitmap is needed during a frequent itemset operation, it is possible to load selected portions of the bitmap into volatile memory, rather than the entire bitmap.

[0112] Various techniques for processing bitstrings that are stored in a bitmap index are described in U.S. Patent Numbers 6,067,540, 6,141,656, 5,899,988, 5,884,307, 5,907,297, 5,963,935, 6,195,656, 6,205,442, and 6,266,662, the contents of which are incorporated herein by this reference. These and other such bitmap handling techniques may be used to increase the efficiency and decrease the resource usage of bitmap indexes that are used to store previously-generated bitmaps for frequent itemset operations.

[0113] By using the clustering techniques described above, the amount of I/O between volatile and non-volatile memory required during a frequent itemset operation is significantly reduced. For example, when the frequent itemset operation begins to evaluate the combinations that belong to a cluster, the database server traverses the bitmap index to locate

the index entry for the appropriate cluster bitmap, and reads the cluster bitmap into volatile memory. The cluster bitmap is then used to evaluate all of the combinations in the cluster, and is then discarded from volatile memory. Thus, by clustering the combinations and using the cluster bitmap, it is not necessary to retrieve a new bitmap from the bitmap index every time a new combination is being evaluated.

PREFIX TREE COUNTING

[0114] In the previous sections, the number of occurrences of a combination is determined by (1) generating a bitmap for the combination, and then (2) counting the TRUE values in the bitmap. Once the number of occurrences of the combination is determined, whether the combination satisfies the frequency criteria may be determined based on the number of occurrences, the total number of item groups in the item group population, and the frequency threshold. The technique of determining the number of occurrences of a combination based on the bitmap for the combination is referred to herein as the bitmap intersection technique.

[0115] Various other techniques may be used to determine the number of occurrences of a combination. For example, an alternative technique is referred to herein as the prefix tree counting technique. When executing the prefix tree counting technique, at the start of a phase, the database server builds a "prefix tree" based on the candidates for that phase. The tree is N levels deep, where N is the number of items in the combinations that are being considered during the phase. Each node in the tree is associated with an item.

[0116] For example, assume that the candidates at the beginning of a three-item phase are:

1. {apple, banana, milk}

2. {apple, banana, orange}
3. {apple, milk, orange}
4. {banana, milk, orange}

[0117] Based on these candidates, the prefix tree 500 in FIG. 5a may be generated. Each leaf node of the tree represents an N-item combination that includes (1) the item associated with the leaf node and (2) the items associated with the nodes that must be traversed to reach the leaf node. For example, leaf node 502 represents the combination (apple, banana milk), while leaf node 504 represents the combination (banana, milk, orange).

[0118] Once the prefix tree has been established, the database server then iterates through each item group in the population, incrementing counters at each leaf node that represents a combination of the item group. For example, assume that the item group population consists of the item groups:

- T1 - (apple, banana, milk)
- T2 - (apple, banana, milk, orange)
- T3 - (orange)
- T4 - (banana, milk, orange)
- T5 - (apple, milk, orange)
- T6 - (apple, banana, orange)

[0119] Item group T1 only contains one three-item combination (apple, banana, milk). Therefore, to process item group T1, the database server traverses tree 500 once based on the combination (apple, banana, milk), and increments a counter associated with the corresponding leaf node (leaf node 502).

[0120] Item group T2 includes several three-item combinations, namely (apple, banana, milk), (apple, banana, orange), (apple, milk, orange), and (banana, milk, orange).

Consequently, to process item group T2, the database server traverses tree 500 four times. At the end of each traversal, the counter associated with the resulting leaf node is incremented.

[0121] After every item group has been processed in this manner, the counters associated with each leaf node will indicate the number of occurrences of the combination associated with the leaf node. In the present example, after processing item groups T1 through T6, the counters of prefix tree 500 will have all been incremented to 2, as illustrated in FIG. 5b. Thus, the three-item combination with each of the leaf nodes occurs twice in the item group population.

DYNAMIC SELECTION OF COUNTING TECHNIQUE

[0122] The bitmap intersection technique and the prefix tree technique are merely two examples of techniques for counting the number of occurrences of the candidate combinations within the item group population. Other techniques may also be employed. However, the efficiency and performance of any given technique can vary widely based on a variety of factors, including the current workload of the computer system, the amount of resources available (e.g. volatile memory), and numerous other factors that affect the I/O cost and CPU cost of the techniques. Thus, while a first technique may perform vastly superior to a second technique under certain conditions, the second technique may perform vastly superior to the first technique under other conditions.

[0123] According to one aspect of the invention, the database server evaluates one or more factors at the start of a phase, and dynamically selects which occurrence counting technique to use during that phase based on those factors. Thus, because the occurrence counting technique is selected on a phase-by-phase basis, it is possible for several different occurrence counting techniques to be used in a single frequent itemset operation. For

example, the bitmap intersection technique may be used during the 1-item, 2-item and 5-item phases of an frequent itemset operation. During the 3-item and 6-item phase of the same frequent itemset operation, the prefix tree technique may be used. During the 4-item phase of the same frequent itemset operation, a technique that differs from both the bitmap intersection technique and the prefix tree technique may be used.

DETERMINING WHICH OCCURRENCE COUNTING TECHNIQUE TO USE

[0124] The determination as to which of several occurrence counting techniques to use during any give phase may take into account a variety of factors. The specific factors, and the weights given to each, may vary from implementation to implementation. The present invention is not limited to any particular occurrence counting techniques, factors, or weights.

[0125] According to one embodiment, at the start of a phase of a frequent itemset operation, the database server selects between the bitmap intersection technique and the prefix tree technique based on both the I/O and CPU costs of each. After estimating the two costs for each technique, the two costs are combined with some weighting factors to generate a combined cost for each technique. Finally, the combined costs for the two techniques are compared with each other, and the technique with the lower cost is selected for the current phase of the frequent itemset operation. An example of how I/O costs and CPU costs can be calculated for each of these techniques is described hereafter. However, the actual cost calculations used may vary from implementation to implementation.

DETERMINING I/O COST FOR THE PREFIX TREE TECHNIQUE

[0126] According to one embodiment, the I/O cost for the prefix tree technique is determined based on the formula:

$$IO = transaction_list_reading_cost * (treeSize / memSize + 1)$$

where treeSize indicates the size of the candidate prefix tree, and

memSize indicates the size of the memory that can be used for prefix tree.

[0127] Because the entire candidate prefix tree may not fit in memory, the prefix tree can be counted in partitions. Under these circumstances, all the transaction lists are read for each partition of the tree.

DETERMINING I/O COST FOR THE BITMAP INTERSECTION TECHNIQUE

[0128] According to one embodiment, the I/O cost for the bitmap intersection technique is determined based on the formula:

$$IO = bitmaps_reading_cost + index_access_cost$$

where index_access_cost = number_index_access * each_bitmap_access_cost,

bitmaps_reading_cost is the cost reading all the frequent items' bitmaps,

number_index_access is number of index accesses to get the bitmap in if the bitmap is not in volatile memory, and

each_bitmap_access_cost is for each frequent item's bitmap reading cost.

[0129] The bitmap intersection technique may involve index access because there is no guarantee that all the bitmaps can fit in volatile memory. In fact, even the intermediate bitmaps may not all fit in volatile memory. Under these circumstances, the intermediate bitmaps can be dumped to non-volatile memory, and then read in to volatile memory when needed.

DETERMINING CPU COST OF PREFIX TREE TECHNIQUE

[0130] According to one embodiment, the CPU cost for the prefix tree technique is determined based on the formula:

$$\text{CPU} = \text{tree_creation_cost} + \\ \text{number_of_transactions} * \text{tree_traversal_cost_per_transaction}$$

where tree_creation_cost is the cost of creating the prefix tree,
number_of_transactions is the total number of item groups, and
tree_traversal_cost_per_transaction is the cost of traversing the tree with a item group.

[0131] The following recursive function is an example of a function by which the tree traversal cost may be estimated:

```
traverse_cost(current_depth, tree_fanout, transaction_length)
{
    if (current_depth = leaf_level_depth) then
    {
        leaf_level_cost := (compare tree_fanout nodes + increment counters)
        return leaf_level_cost;
    }
    else
    {
        intermediate_cost := 0;
        for (i := 1; i < transaction_length - current_depth; i++)
        {
            intermediate_cost := (compare tree_fanout/transaction_length nodes)
```

```

        + traverse_cost(current_depth - 1, tree_fanout,
            transaction_length - i * tree_fanout/transaction_length);
    }
    return intermediate_cost
}
}

```

[0132] In general, the recursive function determines the leaf node cost by comparing number of tree_fanout nodes' cost and the increment counters' cost. The intermediate node cost is determined by the cost of access all the children nodes and the compare node cost in the intermediate node.

[0133] According to one embodiment, an estimate of the traversing cost is determined based on the current_phase, fanout, and average_transaction_length. The cost of traversing an average transaction may be multiplied by the total number of transactions to arrive at a total traversing cost.

DETERMINING THE CPU COST OF THE BITMAP INTERSECTION TECHNIQUE

[0134] According to one embodiment, the CPU cost for the bitmap intersection technique is determined based on the formula:

$$\text{CPU} = \text{number_of_intersection} * \text{intersection_cost_per_bitmap};$$

where number_of_intersection is the number_of_candidates * rolldown factor,

number_of_intersection is the total number of intersections needed in the rolldown model, and

intersection_cost_per_bitmap is the intersection cost of two bitmaps.

GENERATING TOTAL COST VALUES

[0135] According to one embodiment, after both CPU and I/O cost estimates are calculated for each occurrence counting technique, a total cost is generated for each technique. According to one embodiment, the total cost of a technique is determined based on the formula:

$$\text{cmb_cost} = \text{combine_with_weight}(\text{io_cost}, \text{cpu_cost})$$

[0136] After a combined cost has been determined for each occurrence counting technique, the occurrence counting technique with the lowest estimated cost is chosen as the occurrence counting technique to use for the current phase of the frequent itemset operation.

[0137] It should be noted that the specific weights given to each of the constituent costs may vary from implementation to implementation. Further, embodiments that dynamically select an occurrence counting technique on a per-phase basis during a frequent itemset operation are not limited to any specific type of constituent costs, how the constituent costs are determined, nor the relative weights given to the constituent costs.

[0138] It should be noted that each occurrence counting technique will typically have a startup cost. For example, for the bitmap intersection technique, single-item bitmaps have to be prepared. In the prefix tree technique, transaction lists have to be prepared. In one embodiment, unnecessary start up costs are avoided by making an initial determination about whether to always use only one of the available occurrence counting techniques. Such a determination may be based, for example, on the number of candidate combinations that may have to be considered during the frequent itemset operation, and the size of the item group population. If it is determined that only one occurrence counting technique will be used during the frequent itemset operation, then the start up costs for the other occurrence counting techniques are avoided.

[0139] Various other techniques may be used to avoid or delay start up costs. For example, at the beginning of the frequent itemset operation, it may be determined that only a subset of the available occurrence counting techniques will be considered during the frequent itemset operation. Consequently, the start up costs of only the selected subset of occurrence counting techniques are incurred. Further, the preparations necessary to use a particular technique may be delayed until it becomes clear that the technique will actually be considered or used during a phase of the frequent itemset operation.

HARDWARE OVERVIEW

[0140] Figure 6 is a block diagram that illustrates a computer system 600 upon which an embodiment of the invention may be implemented. Computer system 600 includes a bus 602 or other communication mechanism for communicating information, and a processor 604 coupled with bus 602 for processing information. Computer system 600 also includes a main memory 606, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 602 for storing information and instructions to be executed by processor 604. Main memory 606 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 604. Computer system 600 further includes a read only memory (ROM) 608 or other static storage device coupled to bus 602 for storing static information and instructions for processor 604. A storage device 610, such as a magnetic disk or optical disk, is provided and coupled to bus 602 for storing information and instructions.

[0141] Computer system 600 may be coupled via bus 602 to a display 612, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 614, including alphanumeric and other keys, is coupled to bus 602 for communicating information

and command selections to processor 604. Another type of user input device is cursor control 616, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 604 and for controlling cursor movement on display 612. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0142] The invention is related to the use of computer system 600 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 600 in response to processor 604 executing one or more sequences of one or more instructions contained in main memory 606. Such instructions may be read into main memory 606 from another computer-readable medium, such as storage device 610. Execution of the sequences of instructions contained in main memory 606 causes processor 604 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0143] The term “computer-readable medium” as used herein refers to any medium that participates in providing instructions to processor 604 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 610. Volatile media includes dynamic memory, such as main memory 606. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 602. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

[0144] Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0145] Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 604 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 600 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 602. Bus 602 carries the data to main memory 606, from which processor 604 retrieves and executes the instructions. The instructions received by main memory 606 may optionally be stored on storage device 610 either before or after execution by processor 604.

[0146] Computer system 600 also includes a communication interface 618 coupled to bus 602. Communication interface 618 provides a two-way data communication coupling to a network link 620 that is connected to a local network 622. For example, communication interface 618 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 618 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 618 sends and receives

electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0147] Network link 620 typically provides data communication through one or more networks to other data devices. For example, network link 620 may provide a connection through local network 622 to a host computer 624 or to data equipment operated by an Internet Service Provider (ISP) 626. ISP 626 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the “Internet” 628. Local network 622 and Internet 628 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 620 and through communication interface 618, which carry the digital data to and from computer system 600, are exemplary forms of carrier waves transporting the information.

[0148] Computer system 600 can send messages and receive data, including program code, through the network(s), network link 620 and communication interface 618. In the Internet example, a server 630 might transmit a requested code for an application program through Internet 628, ISP 626, local network 622 and communication interface 618.

[0149] The received code may be executed by processor 604 as it is received, and/or stored in storage device 610, or other non-volatile storage for later execution. In this manner, computer system 600 may obtain application code in the form of a carrier wave.

[0150] In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this application, in the specific form in which such claims issue, including any subsequent

correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.
